

IN THE SPECIFICATION:

Please replace the first full paragraph of specification page 1 with the following replacement paragraph:

— RELATED APPLICATIONS

This is a Continuation-In-Part application of U.S. ~~patent application~~Patent Appli-
cation Serial No. 09/625,234 entitled NEGOTIATING TAKEOVER IN HIGH AVAIL-
ABILITY CLUSTER, filed July 25, 2000, now U.S. Patent No. 6,728,897, issued on
April 27, 2004. This application is related to copending U.S. Patent Application Serial
No. 09/933,866 entitled OPERATOR INITIATED GRACEFUL TAKEOVER IN A
NODE CLUSTER, filed August 20, 2001, now issued as U.S. Patent No. _____. —

Please replace the last full paragraph of specification page 6 with the following replacement paragraph:

— The teaching of this invention can be adapted to a variety of storage system architectures including, but not limited to, a network-attached storage environment, a storage area network and disk ~~assembly~~assemblies directly-attached to a client/host computer. The term “storage system” should therefore be taken broadly to include such arrangements. It is expressly contemplated that the various processes, architectures and procedures described herein can be implemented in hardware, firmware or software, consisting of a computer-readable medium including program instructions that perform a series of steps. However, it should be understood that the teaching of this invention can be applied to any server systems. —

Please replace the first full paragraph of specification page 8 with the following replacement paragraph:

— As part of a cluster operation, filers A & B have primarily assigned to each of them a disk shelf 160 comprised of hard disk storage devices D1 – Dn that operate in a manner well known in the art. The filers are controlled by a storage operating system, which may preferably be the Data ONTAP™ storage operating system available from Network Appliance, Inc that is optimized to provide filer services. To understand the failover operation described further in this specification, it is important to understand that filers A & B access both disk shelves A and B. Filer A accesses its disk shelf A via loop A 154, through a transmission medium 157, and accesses disk shelf B via loop B 155, through a transmission medium 156. Similarly, filer B has primarily assigned to it a disk shelf B that it accesses via its loop A 154, through a transmission medium 158, and accesses disk shelf A via its loop B 154, through a transmission medium 159. This joint access is necessary for a partner filer to access a failed filer's disk shelf to continue providing file services to the clients of the failed filer after a takeover. —

Please replace the last full paragraph of specification page 8 with the following replacement paragraph:

— Each filer has a failover monitor 400 that continuously checks and records the status of hardware and software associated with the filer.- This information is kept in NVRAM 151 in each filer. More details of the operation of a failover monitor are described in this specification with reference to Figure 4. Other persistent storage means or a removable storage media may also be used instead of NVRAM. —

Please replace the second full paragraph of specification page 9 with the following replacement paragraph:

— Fig. 2 is a block diagram of filer 200 comprising a processor 202, cluster interconnect 153, NVRAM 151, a memory 204 that may store an application 220, a storage adapter 206 and at least one network adapter 208 all interconnected by a system bus 210, which is preferably a conventional peripheral computer interconnect (PCI) bus. Storage adapter 206 is connected to disks 216 via a Fibre Channel link. The filer also includes ~~the preferable~~ storage operating system 230 stored in memory 204 that implements a file system to logically organize information stored as a hierarchical structure of directories and files on the disks in an assigned disk shelf 212. Disks in the disk shelf are typically organized as a RAID 4 (Redundant Arrays of Inexpensive Disks) array to protect against data loss caused by disk failure in a manner well known in the art. RAID arrays also improve data availability because a filer can continue operation even with a single failed disk. —

Please replace the last full paragraph of specification page 12 with the following replacement paragraph:

— In the event that the failed filer has not shut down at the end of the countdown period, the partner sends a “takeover” command to the failed filer, thereby forcing it to shutdown, asserts disk reservations, takes over responsibility of the disks of the failed filer, and takes over the services of the failed filer. —

Please replace the first full paragraph of specification page 13 with the following replacement paragraph:

— In addition, to prevent the failed filer from coming back online of its own accord, the partner may, in some instances, periodically sends-send a “please die” commands to the failed filer over the cluster interconnect to assure that it remains out of service. —

Please replace the first full paragraph of specification page 14 with the following replacement paragraphs:

— In addition, at block 505 the partner filer investigates the problem, as requested by the failed filer, by attempting to accomplish what the failed filer indicates it cannot do. The goal of the cooperative testing is to confirm a problem with the operation of the failed filer through self-diagnosis or collective intelligence with the assistance of the partner filer, which it does at block 507.

At step 509 the failed filer requests that its partner filer takeover its operations and also indicates to its partner filer the type of problem(s) it has detected. At block 511, before the partner filer takes over the operations of the failed ~~partner filer~~ it first determines if it is able to do so. —

Please replace the last two full paragraphs of specification page 14 with the following replacement paragraphs:

— In the event that the failed filer has not completed its operations and shut itself down by the end of the countdown period, the partner filer forces the failed filer to shut-down at block 519 by issuing a “takeover” command and asserting disk reservations, thereby taking over responsibility of the disks of the failed filer, and providing the file services of the failed filer at block 521.

After detected problems are fixed, which may be as simple as rebooting of the failed filer, the identity and replicated network interfaces and network addresses used by the partner filer are discontinued and control is returned to the restored filer. The program then returns to step 501 where each filer provides file services, and is monitoring its own operations until an operational problem is again detected. —